

**** 2007.09.04**

***** Releases**

No new releases. Client tested by Steve - seems to pass (see [2007.08.14](#) for details, along with SAM_GROUP).

DB server - Randolph's retired date and restricting dataset definition names are pending.

SeedData and DBTestLoader packages were updated for the retired date stuff, but have not been updated to upd. Anil is supposed to do this (This is from the composite index over retired date and file name).

***** D0**

No core sam issues. Quiet week.

Ora2 is down all morning.

d0rsam01 seems to be failing.

SAMGrid problems: IT-2974 (site specific issue)
https://plone3.fnal.gov/SAMGrid/tracking/2974/pcng_issue_view
Condor was not started with the right uid permissions.

d0srv071 is a input cache machine. d0srv072 is not used anymore. 16-18 TB of cache each. Perhaps we can task this machine.

Power outage on Saturday (will stop clued0 SAM station - may be down all weekend if problems). Hopefully
DNS for name service has been corrected for ora2.

***** CDF**

Power outage today.

All nodes below fcdldata093 inclusive should be retired from dCache.

dCache and enstore will be upgraded today since downtime.

***** MINOS**

No report

***** DB**

No report

***** dCache/enstore**

No report

CDF SAM shift report (week 08.28 to 09.03)

1 message

Stefano <stefano.giagu@gmail.com>

Tue, Sep 4, 2007 at 3:05 AM

To: sam-design@fnal.gov, CDF DH Group <cdfdh_oper@fnal.gov>

Hi,

here follows the SAM shift report for the week 08.29 to 09.03.
In general a somehow quiet week, most relevant issues:

Tue. Aug. 28th:

- rawadata_prd died at cdfsamraw1 (fixed) **upgraded by Dehong**
- dcache pool fcdldata138 3,4 went offline, restarted by Alex (fixed)

Wed. Aug. 29th:

- mover 9940B12 offline, mail sent to enstore-admin (fixed)
- still problems with dcache pools fcdldata138 1,2,3,4 offline, nodes un-pingeable, mail sent to dcache-admin, ticket opened (fixed)

Thu. Aug 30th:

- quickly rising number of files to restore, problem related to fcdldata316, Krzysztof moved the node to TestPool, then investigated by Alex: "Pool processes grew up to ~1GB RSS and there was large number of open pipes (one pool had ~700 pipes open)", restarted, and put back to readPool.
- large number of errors in upload server 322 due to a unregistered dataset, mail sent to cdf_mc_ops, dataset registered by Craig (fixed)
- mover LTO3_18 dead, mail sent to enstore-admin, due to troubles with one of the handbots in the SL8500 at GCC (fixed) **2nd time handbot had problems.**

Fri. Aug. 31st:

- (from Alex) Many stage transfers in "Waiting" state, at least two pool shoed "too many open files" errors and there was large number of open pipes, restarted stage pools: fcdldata137-1, fcdldata139-1 and fcdldata154-1 and issued retry to transfer on other pools (fixed)

Sat. Sep. 1st:

- everything OK

Sun. Sep. 2nd:

- again problems with pool fcdldata316, same behavior as two days ago: rapid rise of the number of files to restore with time, fcdldata316 always listed in DH at a glance as the pool with the most **When 316 dies, all clients probably go away!**

file restores and the oldest file restore.

Problem fixed by Timur&Friends (fixed) **Did this from CHEP!**

Mon. Sep 3rd:

- mover LT03_02 dead, mail sent to enstore-admin (still dead)
- enstore encp history monitor reported a bunch of errors like:
"encp version too old: v3_5. Must be not older than v3_6c :
(SRC:/tmp/r2s_arch30457.cry)
(DST:/pnfs/cdf/cdfsys/online/20070901_data/s077.tar.gz.cry)
mail sent to enstore-admin, erros disappeared after ~1 hours (fixed)

Cheers, Stefano
